

**ABSTRACT OF THE DISCLOSURE**

Coordinated SYN denial of service (CSDoS) attacks are reduced or eliminated by a process that instructs a layer 4-7 switch to divert a small fraction of SYN packets destined to a server S to a web guard processor. The web guard processor acts as a termination point in the connection with the one or more clients from which the packets originated, and upon the establishment of a first TCP connection with a legitimate client, opens a new TCP connection to the server and transfers the data between these two connections. It also monitors the number of timed-out connections to each client. When a CSDoS attack is in progress, the number of the forged attack packets and hence the number of timed-out connections increases significantly. If this number exceeds a predetermined threshold amount, the web guard processor declares that this server is under attack. It then reprograms the switch to divert all traffic (i.e. SYN packets) destined to this server to the web guard processor, or to delete all SYN packets to the server in question. If the number of timed-out connections increases, it can also inform other web guard processor arrangements, and/or try to find the real originating hosts for the forged packets. In either event, the server is thus shielded from, and does not feel the effects of, the DoS attack. Alternatively, a simpler approach is to arrange layer 4-7 switches to forward SYN packets to respective "null-cache" TCP proxies that each are arranged to operate without an associated cache, and therefore be inexpensive to install and operate. These null-cache TCP proxies, when subject to a CSDoS attack, will not successfully establish a TCP connection with a malicious host, due to the nature of the attack itself. Accordingly, no connections will be made from the null-cache TCP proxies to the server under attack, and the server will be protected.